

**A Performance History of Several Multitasking
Codes on the NAS Y-MP: 4/15/92 Update**

*Robert J. Bergeron*¹

Report RND-92-019 November 1992

NAS Systems Development Branch
NAS Systems Division
NASA Ames Research Center
Mail Stop 258-6
Moffett Field, CA 94035-1000

¹ Computer Sciences Corporation, NASA Ames Research Center, Moffett Field, CA 94035-1000

A Performance History of Several Multitasking Codes on the NAS Y-MP: 04/15/92 Update

Robert J. Bergeron
Computer Science Corporation
NASA Ames Research Center
Numerical Aerodynamic Simulation Systems Division

ABSTRACT

This paper updates performance histories of codes used to ensure that the Y-MP maintains a sustained multitasking capability across the major Cray upgrades. The history of the ARC3D code shows that this capability should be tested to ensure that upgrades do not impair performance. Measurements indicate that Cray hardware and software upgrades in the 1991-1992 period have maintained the performance of the test multitasking codes.

INTRODUCTION

This paper updates the performance histories of codes used to ensure that the Cray Y-MP maintains a sustained multitasking capability across the major Cray upgrades. Two of the codes, ARC3D and THRED, provided verification that the Y-MP could sustain 1 GFLOP/sec during the NAS acceptance testing of Y-MP Serial Number (SN) 1002. Recent upgrades have motivated the addition of two other codes to this history. The third code, SPARK, provides useful data on the state of the autotasker and a fourth code, MXM, provides information on the efficiency of the operating system with respect to autotasking. The monitoring emphasizes performance on major upgrades, that is, new versions of the operating system and compiler. All results reported herein reflect calculations performed on a dedicated machine.

The major event during the April 1991-April 1992 period was the hardware upgrade from 128 MW to 256 MW. A second event impacting ARC3D and THRED performance was the replacement of the microtasking libraries with the autotasked libraries.

The following tables provide the testing histories; the letter N in the first column denotes new performance data. In addition to the CFT77 compiler version, the table also provides the versions of the Fortran PreProcessor (FPP)

and Fortran MidProcessor. These utilities enable Fortran codes to exploit parallel processing. FPP parses the user's Fortran source to produce a transformed source file containing parallel directives. FMP converts the parallel directives in the transformed source into Cray intrinsic functions or library calls. The CFT77 compiler can then use the FMP-produced code to create a parallel executable code.

ARC3D

The installation of the Session Reservable File System (SRFS) was completed during this period. Because SRFS treats the SSD space as a filesystem, dedicated testing of ARC3D now requires the user to reserve 60 MW of SSD file space before the dedicated testing period, or to perform a SRFS backup before and a restore after the dedicated time.

Since ARC3D is highly parallel and I/O intensive, its dedicated performance will always be sensitive to both the UNICOS operating system environment and any local NAS modifications. Current execution of ARC3D occurs with network daemons activated to allow monitoring from workstations and an increase in interference would be a potential cause of degraded performance.

Figure 1 depicts the performance history for ARC3D. The performance minimum in November of 1989 reflects the UNICOS multitasking problems discussed previously (Bergeron, 1991). The table shows that eight executions of ARC3D occurred during the current period and Figure 1 shows a performance level about 8% below the 1991 peak. The plot data exclude duplicate runs (e.g., 07/30/91) or "check" runs made with old versions of the compiler (e.g., 08/06/91 and 10/09/91). Reasons for the reduced performance include the SRFS implementation, the compiler upgrade and the replacement of the microtasked libraries by the autotasked libraries.

Date	FPP	FMP	Compiler	UNICOS	Seconds	GFLOP/sec	Comments
10/13/88	N/A	N/R	3.0	4.0.6+	1114	1.076	1002 Acceptance Test
06/29/89	N/A	N/R	3.0	5.0.8	1525	0.786	UNICOS problems
11/02/89	N/A	N/R	3.0.1.16	5.0.12	1543	0.777	UNICOS problems
11/08/89	N/A	N/R	3.0.1.16	5.0.12	1100	1.090	Cray test with SDS
11/08/89	N/A	N/R	3.0.1.16	5.0.12	1164	1.030	SSD + new unit size
02/03/90	N/A	N/R	3.0.2.2	5.0.13	1086	1.104	1030 Upgrade
04/17/90	-----	-----	-----	5.1.8	----	-----	5.1 upgrade
11/01/90	-----	-----	-----	6.0.1	----	-----	6.0 upgrade
03/13/91	N/A	4.0.3(9)	4.0.3.1	6.0.9	1147	1.045	6.0 File Structure
04/24/91	N/A	5.0 (29)	4.0.3.1	6.0.11	1036	1.157	NAS test with SDS
N07/30/91	N/A	5.0 (37)	5.0.0.1	6.0.12	1125	1.065	New cft77 and SRFS
N08/01/91	N/A	5.0 (37)	5.0.0.1	6.0.12	1120	1.070	New cft77 and SRFS
N08/06/91	N/A	5.0 (29)	4.0.3.1	6.0.12	1057	1.134	Old cft77 and SRFS
N09/25/91	N/A	5.0 (37)	5.0.0.1	6.1.4	1123	1.066	cft77 5.0
N10/09/91	N/A	4.0.3(9)	4.0.3.1	6.1.5	1060	1.131	Old cft77
N10/09/91	3.03M1	5.0 (37)	5.0.0.1	6.1.5	1110	1.078	Autotasked libraries
N02/13/92	3.03M5	5.0.1(6)	5.0.1.1	6.1.5	1108	1.080	256 MW upgrade,
--	---	---	---	---	--	--	MP_DEDICATED=0
N03/04/92	3.03N6	5.0.1(10)	5.0.1.18	6.1.5	1104	1.084	256 MW upgrade,
--	---	---	---	---	--	--	MP_DEDICATED=1

N/A-Not Applicable

N/R-Not Recorded

The flexibility of the NAS SRFS causes ARC3D I/O to be about 5% slower than the fastest possible Y-MP I/O path, the SDS. A separate program always measures ARC3D I/O before a dedicated run to establish its effect on the ARC3D performance. Comparison of ARC3D performance on 04/24/91 with the result in 08/06/91 indicates the SRFS I/O reduction translates into a 2% performance degradation.

Comparison of the 10/09/91 performance with the 09/25/91 performance indicates that version 5.0 of the compiler is about 6% slower than version 4.0. Examination of the compiler listing revealed that cft77 5.0 had vectorized 3 additional small loops. However, the two object files differ substantially, and the 4.0 version invokes the microtasking library to manage the parallel code. These differences produce the differing performances for the two compilers.

Replacement of the microtasking libraries with the autotasking libraries required that the environmental variable, MP_DEDICATED, be set to 1 to ensure that the extra processors remain with the master process during periods of singletasked activity, e.g., I/O activity on the SSD or other filesystems. Comparison of the 02/13/92 performance with the 03/04/92 performance indicates an improvement of less than 1% change due to MP_DEDICATED for ARC3D.

In this period, the ARC3D performance declined about 8% with 6% of the decline due to cft77 5.0 upgrade and 2% due to the NAS SRFS modification. Execution with autotasked libraries (as will be required in cft77 version 6.0) produced no substantial performance difference.

THRED

Five executions of THRED occurred during the testing period. Although the test version of THRED performs file I/O, the code does not employ the SSD, and the SRFS changes had no impact on its performance.

Comparison of the 03/13/91 result with the 07/25/91 performance shows that the compiler upgrade to version 5.0 made essentially no impact on THRED's performance. Wall-clock times for THRED on a dedicated testing period are reproducible to the nearest second, but a rerun of the test at a later date with all known conditions the same can result in a few (plus or minus 10) seconds difference.

Date	FPP	FMP	Compiler	UNICOS	Seconds	GFLOP/sec	Comments
10/13/88	N/A	N/R	3.0	4.0.6+	1218	1.227	1002 Acceptance
06/29/89	N/A	N/R	3.0	5.0.8	1190	1.255	Minor clock problem
02/03/90	N/A	N/R	3.0.2.2	5.0.13	1080	1.383	1030 Upgrade
04/17/90	-----	-----	-----	5.1.8	----	-----	5.1 upgrade
11/01/90	-----	-----	-----	6.0.1	----	-----	6.0 upgrade
03/13/91	N/A	4.0.3(9)	4.0.2.1	6.0.9	1057	1.413	6.0 File Structure
N07/25/91	N/A	5.0 (31)	5.0.0.1	6.0.12	1061	1.408	New cft77
N09/25/91	N/A	5.0 (37)	5.0.0.1	6.1.4	1049	1.425	UNICOS upgrade
N10/09/91	3.03M1	5.0 (37)	5.0.0.1	6.1.5	1069	1.398	Autotasked libraries
--	---	---	---	---	--	--	MP_DEDICATED=0
N02/05/92	3.03M5	5.0.1(6)	5.0.1.1	6.1.5	1102	1.356	256 MW upgrade
--	---	---	---	---	--	--	MP_DEDICATED=0
N02/26/92	3.03N6	5.0.1(10)	5.0.1.18	6.1.5	1056	1.415	MP_DEDICATED=1

N/A-Not Applicable

N/R-Not Recorded

Replacement of the microtasking libraries with the autotasking libraries means that an environmental variable, MP_DEDICATED, controls the scheduling of microtasked jobs in dedicated time. Setting MP_DEDICATED to 1 ensures that the extra processors remain with the master process during periods of singletasked activity. The two THRED (10/09/91 and 02/05/92) tests which executed with this variable improperly set to 0 displayed depressed performance rates. Execution with MP_DEDICATED properly set gave a value of 1.415 GFLOPS which is the same performance as was measured a year ago. Figure 1 depicts the performance history for THRED. The performance tends to drift gradually upward as Cray upgrades its compilers and operating systems. The two reductions in THRED performance which occurred in this period resulted from the suboptimal setting of the MP_DEDICATED environment variable.

SPARK

Three dedicated executions of the SPARK code occurred during this period and Figure 1 shows a 9% performance improvement during the year. The number of autotasked constructs decreased slightly but a decrease in the run-time scheduling values allowed more loops to be executed in parallel. The table contains results for two sizes of the code, the smaller size executing 12.73 billion floating point operations, and a larger version executing 110.67 billion floating point operations.

Date	FPP	FMP	Compiler	UNICOS	Seconds	GFLOP/sec	Comments
5.1 Option							
12/14/89	2.26B18	3.1 (33)	3.0.2.2	5.0.13	18.18	0.700	12.73 GFLOPS
12/14/89	2.24S6	-----	3.0.2.2	5.0.13	20.39	0.624	
12/14/89	2.24S6	-----	3.0.2.2	5.0.13	15.37	0.828	Old Libraries
07/13/90	2.24S6	-----	3.0.2.2	5.1.10	14.49	0.879	
07/13/90	3.00Z51	4.0.2(12)	4.0.	5.1.10	18.48	0.689	
11/01/90	-----	-----	-----	6.0.1	-----	-----	6.0 upgrade
11/06/90	3.00Z36	4.0.1(38)	4.0.1	6.0.11	16.01	0.795	
04/25/91	3.00Z61	4.0.3(9)	4.0.3.1	6.0.11	151.1	0.732	110.67 GFLOPS
04/25/91	3.03Y4	5.0 (29)	5X402417	6.0.11	142.2	0.778	
N10/23/91	3.03M5	5.0.1(6)	5.0.0.0	6.1.5	142.4	0.777	UNICOS upgrade
N02/13/92	3.03M5	5.0.1(6)	5.0.0.1	6.1.5	130.7	0.846	256 MW upgrade
N03/04/92	3.03N6	5.0.1(10)	5.0.0.18	6.1.5	132.3	0.836	Compiler Upgrade
6.0 Option							
04/25/91	3.00Z61	4.0.3(9)	4.0.3.1	6.0.11	150.8	0.734	
04/25/91	3.03Y4	5.0 (29)	5X402417	6.0.11	143.7	0.770	

MXM

The compiler upgrade to version 5.0 of CFT77 includes a default which allows the autotasker to (silently) replace the autotasked matrix multiply with the SCILIB call SGEMMX. Since MXM is always compiled with default autotasking options, such replacement frustrates the purpose of the test. While an FPP option will disable this default, it is good practice to check the output of the preprocessor before executing this test.

Three tests were run during the the April 1991-April 1992 period. The test on 02/05/92 required 13% more CPU time for successful execution of the vector version and the singletasked version on one processor. The autotasked performance showed a 1% decrease in performance and incorrect efficiencies. Examination of system accounting and wttmp files indicated no unauthorized logins during this time. However, Cray replaced CPU A on 02/08/92 due to its failure to execute a test problem correctly and it seems likely that the 02/05/92 testing may have found an early symptom of the bad CPU.

Date	FPP	FMP		Compiler	UNICOS	Seconds	Gflop/sec	Comments
01/30/91	2.26B18	3.1	(33)	3.0.2.2	6.0.9	0.119	2.278	n=512
01/30/91	2.24S6	-----		3.0.2.2	6.0.9	59.79	2.302	n=4096
02/13/91	3.00Z61	4.0.3	(9)	4.0.3.1	6.0.11	59.79	2.302	n=4096
04/28/91	3.03Y4	5.0	(29)	5X402417	6.0.11	59.91	2.297	n=4096
N09/25/91	3.03M1	5.0	(29)	5.0.0.1	6.1.4	59.91	2.297	n=4096
N02/05/92	3.03M1	5.0.1(6)	5.0.1.1	6.1.5	60.24	2.284	CPU error
N02/13/92	3.03M1	5.0.1(6)	5.0.1.1	6.1.5	59.93	2.296	256 MW upgrade

Efficiency vs NCPUS									
Date	NCPUS								
	1	2	3	4	5	6	7	8	Comments
01/30/91	1.000	0.999	0.992	0.991	0.991	0.988	0.983	0.989	E=0.969 for 16 CPUs
02/13/91	1.000	0.994	0.994	0.997	0.996	0.992	0.991	0.987	E=0.978 for 16 CPUs
N09/25/91	1.000	0.990	0.986	0.987	0.986	0.986	0.984	0.981	E=0.965 for 16 CPUs
N02/05/92	1.000	1.002	1.116	1.116	1.113	1.113	1.110	1.101	CPU error
N02/13/92	1.000	0.993	1.008	1.005	1.005	1.003	1.002	0.999	CPU release error

The test on 02/13/92 displays a performance rate similar to the 09/25/92 measurement. The high efficiencies measured on this date arise from an erroneously large execution time for the MXM singletasked code. The same error occurred for the autotasked code executing on one CPU and the explanation seems to be a bad CPU.

The table shows that the Cray upgrades maintained the performance of autotasked matrix multiply during this period and Figure 1 shows that performance of MXM has been constant throughout testing history.

CONCLUSIONS

Measurements indicate that Cray upgrades during the 1991-1992 period have maintained or increased the performance of three of the test multitasking codes. The ARC3D code displayed about a 6% decrease in performance due to the cft77 5.0 compiler upgrade. The performance of the test codes in dedicated time has remained constant or gradually trended upward.

References

Bergeron, R. J., 1991. A Performance History of Several Multitasking Codes on the NAS Y-MP. NAS Report RND-91-008, July, 1991.

Figure 1



